

CERTIFICATE OF TRANSMISSION

I hereby certify that this correspondence is being transmitted by either electronic submission using the EFS WEB submission system, fax to the U.S. Patent and Trademark office to fax number 571-273-8300, or is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 to on October 9, 2006.

___/David J McKenzie/_____
Attorney for Appellants

PATENT
Docket No. DE920000124US1

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Appellants:	Christian Bolik, et al.)	
)	
Serial No.:	10/015,825)	
)	Group Art
Filed:	December 10, 2001)	Unit: 2145
)	
For:	Method and System for Scalable, High)	
	Performance Hierarchical Storage)	
)	
Examiner:	Ajay M. Bhatia)	
)	

APPELLANTS' REPLY BRIEF

Mail Stop Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

On January 19, 2005, the Appellants filed a timely Notice of Appeal from the Final Office Action mailed September 20, 2005. Along with the Notice of Appeal, the Appellants filed a Petition for Extension of Time and a fee for a one month extension. The Appellants appeal from the rejection of all pending claims.

This Reply Brief is being filed under the provisions of 37 C.F.R. § 41.41 and is in response to the Examiner's Answer mailed on August 9, 2006.

1. REAL PARTY IN INTEREST

The real party in interest is the assignee, International Business Machines Corporation, Armonk, New York.

2. RELATED APPEALS AND INTERFERENCES

There are no related appeals and interferences.

3. STATUS OF CLAIMS

The Examiner acknowledges that the statement of the status of the claims contained in the Appeal Brief is correct.

4. STATUS OF AMENDMENTS

The Examiner acknowledges that the statement of the status of the amendments contained in the Appeal Brief is correct.

5. SUMMARY OF CLAIMED SUBJECT MATTER

The Examiner acknowledges that the statement of the summary of claimed subject matter contained in the Appeal Brief is correct.

6. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

The Examiner acknowledges that the statement of the grounds of rejection to be reviewed on appeal contained in the Appeal Brief is correct.

7. RESPONSE TO EXAMINER'S ARGUMENTS

In response to the Appellants' argument that Cabrera does not teach a method of limiting the scanning scope to less than scanning the entire local storage tree to find candidate files, the Examiner answers that the "limitation the appellant is addressing is the pre-specifying the (sic) scanning scope determined by a number of candidate data files" and that the "specification fails to define how the certain number is specified." Examiner's Answer, mailed 8/9/2006 (hereinafter "Answer") at p. 8. The Examiner then states that "in view of the specification not particularly specifying the "certain number" examiner has interpreted in light of the specification that the application is able to specify the certain number, which Cabrera teaches" and cites column 10, lines 25-30 of Cabrera as evidence. *Id.* The Examiner states that the cited text defines the "certain number" as the group that is selected prior to the migration and that the "claim limitations presented by the applicant are directed to point (sic) at which the (sic) files are migrated, which would (sic) the scanning of files that takes place after the pre-migration, where the applicant has defined the "certain number" a (sic) those that have been identified in the pre-migration." *Id.* at pp. 8-9. The Appellants continue to disagree.

The Examiner's statement that the specification does not particularly specify the "certain number" is incorrect. The specification provides clear support for a first step of "prespecifying a scanning scope determined by a number of candidate data files" followed by "scanning the managed file system until having reached the prespecified number of migration candidate data files" as recited in Claim 1. The certain number of "migration candidate files" is clearly taught so that one of skill in the art understands that the certain number is defined before any scanning

is initiated. The Object and Brief Summary of the Invention clearly states that the “particular step of scanning the managed file system only until having detected **a prespecified amount** of migration candidate files advantageously enables that migration candidates are made available sooner to the migration process wherein migration can be performed as an automigration process not requiring any operator or user interaction.” Specification at p. 5, ll. 16-20 (emphasis added).

The Detailed Description of the Preferred Embodiments clarifies this point and makes it clear that the “amount of files” is equivalent to “the number of files”:

The flow diagram depicted in Fig. 3 illustrates the basic mechanism of managing an HSM system according to the invention. In step 200, **an amount of files, e.g. the number of files** or the entire size of multiple files, for which a scan in the file system shall be performed, is pre-specified. Based on that pre-specified amount, at least part of the file system is scanned 201. It is an important aspect of the invention that not a whole file system is scanned through but only a part of it determined by the pre-specified amount. *Id.* at p. 9, ll. 11-16; *see also id.* at Fig. 3 (emphasis added).

The text of step 200 of Figure 3 also clarifies that the claimed invention limits scanning to less than the entire file system: “Pre-specifying an amount of files to be scanned in one shot.” *Id.* at Fig. 3. The text of step 201 of Figure 3 again makes scanning dependent upon the pre-specified amount: “Serially scanning at least part of the file system, dependent on pre-specified amount.” *Id.* Figures 6a and 6b provide further support for pre-specifying a scanning scope determined by a number of candidate files and then scanning only until reaching the number of candidate files. The specification states that the process scans for “a new bunch of candidates . . . defined by another value MAXCANDIDATES 612 that defines a number of required candidates following candidates criteria.” *Id.* at p. 12, ll. 1-3; *see also id.* at Fig. 6. The specification also states: “Thus the scout process traverses the managed file system in order to find eligible candidates for automigration. Rather than traversing the complete file system it stops as soon as

MAXCANDIDATES 612 eligible candidates were found.” *Id.* at p. 12, ll. 5-8; *see also id.* at Fig. 6. The specification also states: “Instead of attempting to find the “best” migration candidates in one shot, the file system is scanned only until a certain number of migration candidates have been found.” *Id.* at p. 13, ll. 3-4. The specification makes clear to one of ordinary skill in the art that the “certain number” of files is defined through user input, contained in a configuration file, generated by an application, etc. *prior* to scanning.

The citation of the Examiner in support of the Examiner’s argument that Cabrera anticipates the “certain number,” (column 10, lines 25-30) only further makes the Appellants’ point that Cabrera does not limit the scanning scope, but instead searches an entire file system for candidate files. Column 10, lines 25-30 states: “Referring next to FIG. 3, a timeline for migrating data according to the present invention is presented As indicated by FIG. 3, prior to the time that the migration criteria is met (the migration time), the present invention identifies candidates for ‘pre-migration.’ This is illustrated in FIG. 3 by step 64.” Cabrera at col. 10, ll. 25-30. Figure 3 starts with step 64 which is labeled “IDENTIFY PRE-MIGRATED CANDIDATES.” *Id.* at Fig. 3.

Cabrera discloses pre-migrating files that normally would not be migrated under a traditional migration policy. *Id.* at Abstract, col. 10, l. 30 to col. 11, l. 6. Cabrera discloses that ‘pre-migration’ of files is *copying a file to remote file storage but maintaining the source file in place in the local storage rather than replacing it with a stub as in the traditional migration process.* *Id.* at col. 11, ll. 25-46. The only difference between pre-migrating and migrating is that the file in the local file system is left in place for a time rather than being replaced immediately with a stub. Both processes in Cabrera use the same scanning process that traverses

an entire file tree to find candidate files. Cabrera discloses scanning for candidate files for pre-migration by selecting files that would meet the migration policy in every way except the age of the file on the local storage. *Id.* at col. 10, ll. 30-44. Cabrera does not teach a method of limiting the scanning scope to less than scanning the entire local storage tree to find migration candidate files.

The Examiner states that “[t]he claim limitations presented by the applicant are directed to point (sic) at which the files are migrated, which would (sic) the scanning of files that takes place after the pre-migration, wherein the applicant has defined the ‘certain number’ a (sic) those that have been identified in the pre-migration.” Answer at p. 9. The Appellants disagree.

The Examiner seems to suggest that Cabrera’s teaching that coming up with a list of pre-migration candidates using a migration policy just prior to migration is equivalent to the limitations of Claim 1 which includes the steps of pre-specifying a scanning scope determined by a certain number, and then scanning only until reaching the pre-specified number of migration candidate files *in addition* to scanning using a migration policy to find the migration candidate files. Both Cabrera and Claim 1 use a migration policy and end up with migration candidates just before migration, but Cabrera scans the entire file system for pre-migration candidates while Claim 1 limits scanning to a point where a certain number of migration candidate files is reached. Claim 1 limits searching for migration candidates to less than the entire file system. Cabrera searches an entire file system for pre-migration candidates. The distinction is clear. The Appellants assert that the Examiner has failed to make out a *prima facie* case that Cabrera anticipates Claim 1. The Appellants assert that Claim 1 is in condition for allowance.

In response to the Appellants’ argument that Cabrera “does not indicate any method for

scanning the local storage to find the candidate files,” *see* Appeal of Applicants filed 5/19/2006 (hereafter “Appeal”), the Examiner argues that Cabrera anticipates the limitation of “selecting migration candidate data files according to at least one criteria.” Answer at p.9. The Appellants do not disagree that this limitation is taught by Cabrera. This limitation is directed to scanning for migration candidate files using a *migration policy*, such as file size or file age. Specification at p. 9, ll. 17-18. While both Cabrera and Claim 1 teach scanning using a *migration policy*, Cabrera does not teach, disclose, or suggest scanning **until** having reached a pre-specified number of migration candidate files.

The Appellants’ argument in the Appeal Brief makes this point by reviewing each citation in the Examiner’s Office Action mailed 1/9/2006 (hereafter “Office Action”). The Office Action rejects Claim 1 by merely reciting it verbatim and then providing a few citations in Cabrera as evidence of anticipation without any comment. *See* Office Action at pp. 2-3. The Appeal Brief discusses each citation, including column 9, line 53 to column 10, line 17, and points out that none of the citations provided by the Examiner teaches the limitation of scanning **until** having reached a pre-specified number of migration candidate files recited in Claim 1. Specifically, the citation of column 9, line 53 to column 10, line 17 of Cabrera merely restates the traditional HSM process, including scanning using a migration policy, and does not limit scanning to less than the entire file tree that is searched.

The Examiner is not free to pick and choose elements of the claim, but is required to provide evidence of anticipation of each and every claim limitation. *Lindemenn Maschinenfabrik GmbH vs. American Hoist and Derick Co.*, 730 F2d 1452, 221 USPQ 481, 485 (Fed. Cir. 1984) (For a *prima facie* case of anticipation, each and every element of the claimed invention must be

identically disclosed in a single prior art reference; and those elements must be arranged or connected together in a single reference in the same way as specified in the patent claim). As stated above, and in the Appeal Brief, Cabrera does not teach, disclose, or suggest “prespecifying a scanning scope determined by a number of candidate data files” or “scanning the managed file system **until** having reached the **prespecified number** of migration candidate data files.”

The Examiner states that: “Additionally applicant argues that Cabrera does not describe scanning process to find pre-migration files but implies only some sort of scanning. Here applicant has again misinterpreted the application of prior art to appellant’s claims. Appellant’s claims are related to scanning scope, which occurs when appellant’s invention is completing the migration process. Cabrera teaches, (sic) the claims as appellant has claimed in the present invention. Appellant’s (sic) has elected to argue part of Cabrera and not the reference as a whole. Cabrera when reads (sic) a (sic) reference in it (sic) entirety anticipates the presently claimed invention.” Answer at pp. 9-10. The Appellants disagree.

Cabrera teaches traditional scanning of all files, Cabrera at col. 9, ll. 58-60, but not “scanning the managed file system **until** having reached the **prespecified number** of migration candidate data files.” The claimed invention does not merely claim a “scanning process,” as stated by the Examiner, but specifically limits a scan to a pre-specified number of migration candidate files. The Examiner has not provided any evidence that Cabrera teaches this limitation. The Examiner is over generalizing the limitations of Claim 1 to mean *any* migration criteria used to limit the number of migration candidate files. The invention of Claim 1 stops scanning after reaching a pre-specified number of migration files. The invention of Claim 1 uses a migration policy of at least one attribute to scan until the pre-specified number of migration

candidate files is reached. Cabrera teaches scanning an entire file system using a migration policy to find pre-migration candidate files. The general migration policy of Cabrera does not include stopping scanning when a certain number of pre-migration candidate files is reached. Thus Cabrera does not anticipate Claim 1 or the other independent claims.

The Examiner states that “Cabrera teaches, (sic) the claims as appellant has claimed in the present invention” but the Examiner fails to identify in Cabrera where scanning is stopped after reaching a certain number of migration candidate files. The Examiner states that “Appellant’s (sic) has elected to argue part of Cabrera and not the reference as a whole. Cabrera when reads (sic) a (sic) reference in it (sic) entirety anticipates the presently claimed invention.” The burden of presenting a *prima facie* case for anticipation lies with the Examiner, not the Appellants. The Appellants have discussed each and every citation that the Examiner has put forth as evidence that Cabrera anticipates Claim 1 and have shown that the citations do not anticipate each and every claim limitation. The Examiner has had numerous chances to point out in Cabrera where these limitations are taught and has failed to do so. In addition, Cabrera when read as a whole does not anticipate the limitations of Claim 1. The Appellants again assert Claim 1 and the other independent claims are in condition for allowance.

The Examiner states: “Appellant argues that examiner has not given any weight to the amendment and examiner admitted that applicant have (sic) overcome the rejection. Appellant has misunderstood the examiner. Appellant argument, cancelling and moving claim 2 in to (sic) independent claim 1, in addition amending (sic) similar claim limitation (sic) to the other independent claims. Examiner stated in the advisory action that a new rejection addressing the newly amended claims would have to be address (sic), but they did not present new issues since

they were taken from a dependent claim that presently addressed (sic) by the previous rejection.”

Answer at p. 10.

The Examiner fails to address the fact that the Examiner admitted that the Applicants had overcome the prior art. Page 2 of the Advisory Action clearly has a check in the box next to item number 5. Advisory Action mailed 12/21/2005 (hereinafter “Advisory Action”) at p. 2. Item number 5 states: “Applicant’s reply has overcome the following rejection(s): 35 U.S.C. 102(e) Cabrera claims 1, 3-20. The Appellants assert that this admission by the Examiner clearly puts the claims in condition for allowance.

The Applicants’ remarks in the Office Action Response filed 11/21/05 (“Office Action Response”) made specific arguments for the independent claims, which were amended to include the limitations of claim 2, to overcome the prior art of Cabrera. The arguments were not made in previous office action responses and point out that Cabrera does not include the limitations of scanning **until** reaching a **pre-specified number** of migration candidate files found in claim 2. The Examiner did not present a *prima facie* case for the rejection of claim 2 in the Office Actions mailed 1/26/05 and 9/20/05 and the Appellants deserve more than a terse statement in the Advisory Action upholding the rejection. The Appellants maintain that the Examiner did not give any weight to the arguments presented in the Office Action Response and that the claims should have been allowed.

SUMMARY

In view of the foregoing, each of the claims on appeal has been improperly rejected. Reversal of the Examiner's rejection and allowance of the pending claims 1, 3-20 is respectfully requested.

Respectfully submitted,

_____/David J McKenzie/_____

David J. McKenzie
Reg. No. 46,919
Attorney for Appellants

Date: October 9, 2006
8 East Broadway, Suite 600
Salt Lake City, UT 84111
Telephone (801) 994-4646
Fax (801) 322-1054

8. Claims Appendix

Claims involved in the appeal

1. A method of managing a hierarchical storage management (HSM) environment, the environment including at least one HSM server and at least one file server having stored a managed file system, wherein the at least one HSM server and the at least one file server are interconnected via a network and wherein digital data files are migrated temporarily from the at least one file server to the at least one HSM server, the method comprising:
 - providing at least one list for identifying candidate data files to be migrated;
 - prespecifying a scanning scope determined by a number of candidate data files;
 - scanning the managed file system until having reached the prespecified number of migration candidate data files;
 - selecting migration candidate data files according to at least one attribute;
 - recording the selected migration candidate data files in the provided at least one list for identifying candidate data files; and
 - migrating at least part of the selected candidate data files identified in the at least one list for identifying candidate data files from the file server to the HSM server.
2. Canceled
3. The method according to claim 1, wherein the scanning scope is determined by the total amount of data for the candidate data files and wherein the managed file system is

scanned until having the prespecified amount of data.

4. The method according to claim 1, wherein the scanning of the managed file system is resumed at a location of the managed file system where a previous scanning is left off, and continued accordingly.
5. The method according to claim 1, wherein replacing a migrated data file in the managed file system by a stub file providing at least information about the location of the migrated data file on the HSM server.
6. The method according to claim 1, further comprising monitoring a current state of the managed file system and initiating automigration dependent on the monitored current state of the managed file system.
7. The method according to claim 6, comprising the further steps of automigrating candidate data files with respect to the list for identifying candidate data files and assigning a unique identifier to each of the migrated candidate data files.
8. The method according to claim 7, wherein the unique identifier is specific to the underlying file system allowing direct access to a migrated data file.

9. The method according to any of claim 6, wherein providing two lists for identifying candidate data files, whereby the first list is generated and/or updated by a scanning process and whereby the second list is used by a automigration process, and whereby the automigration process gathers the first list from the scanning process when all candidate data files of the second list are migrated.
10. The method according to any of claim 9, wherein the automigration process is performed by a master/slave concept where the master controls the automigration process and selects at least one slave to migrate candidate data files provided by the master.
11. The method according to claim 1, comprising the additional steps of ranking and sorting the candidate data files contained in the at least one list for identifying candidate data files, in particular with respect to the a file size and/or time stamp of the data files contained in the at least one list for identifying candidate data files.
12. The method according to claim 1, wherein the scanning of the managed file system is initiated dependent on expiration of a prespecified wait interval or initiated by the automigration process.

13. A method of reconciling a managed file system migrated from a file server to an hierarchical storage management (HSM) server via a network in accordance with the method according to any of claims 7 to 12, with a current state of the managed file system on the file server, wherein data files migrated to the HSM server are recorded in a list of migrated data files having a unique identifier for each of the migrated data files, the method comprising the steps of:
- querying the list of migrated data files migrated from the managed file server to the HSM server;
 - for each file entry in the list of migrated data files, retrieving from the managed file system at least one attribute of the migrated data file that is identified by the corresponding unique identifier;
 - comparing the retrieved attributes with the corresponding attributes stored in the list of migrated data files; and
 - updating the HSM server for the migrated managed file system dependent on the results of the preceding step of comparing.
14. The method according to claim 13, wherein performing the steps of claim 13 by a reconciling process and wherein the reconciling process requests the list of migrated data files via the network from the HSM server.

15. A hierarchical storage management (HSM) system including at least one HSM server and at least one file server having stored a managed file system, the at least one HSM server and the at least one file server being interconnected via a network, where data files are migrated temporarily from the at least one file server to the at least one HSM, the system comprising:
- means for prespecifying a scanning scope determined by a number of candidate data files
 - means for scanning the file system until having reached the prespecified number of migration candidate files and for identifying candidate data files to be migrated;
 - means for monitoring the managed file system;
 - means for migrating candidate data files to the HSM server;
 - means for reconciling the managed file system.
16. The system according to claim 15, further comprising a means for replacing a migrated data file in the managed file system by a stub file providing at least information about the location of the migrated data file on the HSM server.
17. The system according to claim 15, further comprising means for assigning a unique identifier to at least part of the candidate data files stored in the storage means.

18. The system according to claim 15, further comprising at least two storage means for identifying candidate data files, where the first storage means is generated and/or updated by a scanning process and where the at least second storage means is used by an automigration process, and where the automigration process gathers the content of the first storage means from the scanning process when all candidate data files of the at least second storage means are migrated.
19. A data processing program for execution in a data processing system comprising software code portions for performing a method comprising:
- providing at least one list for identifying candidate data files to be migrated;
 - prespecifying a scanning scope determined by the number of candidate data files;
 - scanning the managed file system until having reached the prespecified number of migration candidate data files;
 - selecting migration candidate data files according to at least one attribute;
 - recording the selected migration candidate data files in the provided at least one list for identifying candidate data files; and
 - migrating at least part of the selected candidate data files identified in the at least one list for identifying candidate data files from the file server to the HSM server.

20. An article of manufacture comprising a program storage medium readable by a processor and embodying one or more instructions executable by the processor to perform a method comprising:
- providing at least one list for identifying candidate data files to be migrated;
 - prespecifying a scanning scope determined by a number of candidate data files;
 - scanning the managed file system until having reached the prespecified number of migration candidate data files;
 - selecting migration candidate data files according to at least one attribute;
 - recording the selected migration candidate data files in the provided at least one list for identifying candidate data files; and
 - migrating at least part of the selected candidate data files identified in the at least one list for identifying candidate data files from the file server to the HSM server.

9. Evidence Appendix

None

10. Related Proceedings Appendix

None